

How a well-adapted immune system is organized

Andreas Mayer,¹ Vijay Balasubramanian,² Thierry Mora,³ and Aleksandra M. Walczak¹

¹*Laboratoire de physique théorique, CNRS, UPMC and École normale supérieure, 24, rue Lhomond, 75005 Paris, France*

²*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104;
and Initiative for the Theoretical Sciences, The Graduate Center,
The City University of New York, NY 10016, USA*

³*Laboratoire de physique statistique, CNRS, UPMC and École normale supérieure, 24, rue Lhomond, 75005 Paris, France*

(Dated: July 28, 2014)

The repertoire of lymphocyte receptors in the adaptive immune system protects organisms from diverse pathogens. A well-adapted repertoire should be tuned to the pathogenic environment to reduce the cost of infections. We develop a general framework for predicting the optimal repertoire that minimizes the cost of infections contracted from a given distribution of pathogens. The theory predicts that the immune system will have more receptors for rare antigens than expected from the frequency of encounters; individuals exposed to the same infections will have sparse repertoires that are largely different, but nevertheless exploit cross-reactivity to provide the same coverage of antigens; and the optimal repertoires can be reached via the dynamics of competitive binding of antigens by receptors, and selective amplification of stimulated receptors. Our results follow from a tension between the statistics of pathogen detection, which favor a broader receptor distribution, and the effects of cross-reactivity, which tend to concentrate the optimal repertoire onto a few highly abundant clones. Our predictions can be tested in high throughput surveys of receptor and pathogen diversity.

The adaptive immune system protects organisms from a great variety of pathogens by maintaining a population of specialized cells, each specific to particular challenges. Together these cells cover the array of potential threats. To recognize pathogens, the immune system relies on receptor proteins expressed on the surface of its main constituents, the B and T lymphocytes. These receptors interact with antigens (small molecular elements making up pathogens), recognize them through specific binding, and initiate the immune response. Each lymphocyte expresses a unique receptor formed from random combinations encoded in the genome. The receptors later undergo selection through the death and division of the lymphocytes that express them, as well as mutations in the case of B lymphocytes. The diversity of the receptor repertoire determines the range of threats that the adaptive immune system can target.

The detailed *composition* of the immune receptor repertoire, and not just its breadth, is important for conferring effective protection against infections. Broadly speaking, a diverse population of receptors will confer wider immunity, and a larger clonal population of a particular receptor will confer more effective immunity against the pathogens to which it is specific. However, there is a tradeoff between diversity and clone sizes because the number of receptors is limited. By selectively proliferating some receptors at the expense of others, the immune system retains a memory of past infections [1], facilitating subsequent immune responses. Furthermore, while infections increase the populations of receptors with the greatest specificity, they can also lead to a reorganization of the immune repertoire as a whole [2].

How should the repertoire be organized to minimize the cost of infections? We develop a framework for an-

swering this question by abstracting key general features of the immune system: the receptor repertoire is bounded in size, receptors are “cross-reactive” (each antigen binds many receptors; each receptor binds many antigens), and the cost of an infection increases with time. Given these general assumptions, we consider a simplified landscape of pathogens, where infections are drawn from a fixed distribution. By simplifying the setting in this way, and independently of the detailed dynamics of immune responses, we arrive at broad insights about the composition of immune repertoires that are optimal for their pathogenic environments.

The theory predicts, counter-intuitively, that the number of receptors specific to rare pathogens will be amplified relative to the probability of encounter, at the expense of receptors for common infections. We also find that two organisms responding to a pathogen distribution will display unique populations of immune receptors, even though their coverage of pathogens will be similar. How can the immune system achieve these sorts of optima? Surprisingly, we find that simple competition between receptor clones can drive the population to the optimal composition for minimizing the cost of infections.

New high throughput methods are making it possible to survey B-cell and T-cell receptor diversity in fish [3, 4], in mice [2, 5] and humans [6–9]. As methods are developed to better characterize pathogenic landscapes and receptor cross-reactivity, predictions for the composition of optimal repertoires derived from our framework can be directly compared with experiments. To arrive at our results we ask how the immune system should be organized to perform its function well, rather than starting with the detailed dynamics of its components. We are proposing that the universal features of the adaptive immune sys-

tem follow simply from general statistical considerations, while the detailed dynamical implementation arises from the historical contingencies of evolution.

I. DEFINITION OF THE PROBLEM

To find the optimal repertoire distribution we must consider the nature of antigen-receptor interactions, and a penalty that the immune system pays for not recognizing antigens. This penalty must reflect the facts that recognition should happen within reasonable time, before the pathogen colony can significantly increase its size; the interactions between the immune receptors and antigen are probabilistic; and not all antigens are equally frequent. We assume that, although the immune system cannot predict precisely which antigens it will encounter and when, it incorporates an estimate of the *probabilities* of their occurrences. We also take these probabilities to be constant in time. This is an idealization grounded in a separation of timescales, which assumes the distribution of antigens remains constant on timescales on which the immune system adapts.

We call Q_a the probability that the next infection will be caused by antigen a (Fig. 1) and model the immune repertoire by a distribution of receptors P_r , from which lymphocytes with the corresponding receptor are drawn at random. During its time in the periphery, an antigen a will encounter and possibly interact with receptors at a rate $\lambda_a(t)$ which increases with time as the pathogen population grows. Each encounter will occur with a different receptor r drawn from P_r . The mean number of encounters between antigens and receptors after a time t , which we will call *effective time*, is defined as $m_a(t) = \int_0^t d\tau \lambda_a(\tau)$, where $t = 0$ is set by the introduction of the antigen. We shall see that the cost of an infection is most easily expressed in terms of the expected number of encounters before recognition, and hence in terms of $m_a(t)$.

An antigen a and a receptor r interact with a certain strength set by the binding affinity between the two molecules. This is described by the probability $f_{r,a}$ that an antigen a colliding with the receptor r results in a recognition event, leading to the activation of the lymphocyte expressing that receptor. $f_{r,a}$ will be called the cross-reactivity function. Each encounter with a random antigen has a probability $\bar{P}_a = \sum_r f_{r,a} P_r$ to lead to recognition and trigger an immune response. Since recognition is a stochastic event, the time t to the first recognition event, or response time, is random and distributed according to the probability distribution function $H_a(t) = \lambda_a(t) \bar{P}_a e^{-m_a(t) \bar{P}_a}$ (see App. A for a derivation).

The longer the system fails to detect the antigen, the more likely the infection is to become harmful. We assume that the integrated harm caused by an antigen since the beginning of an infection is an increasing function $F_a(t)$ of the time of first recognition. The mean harm




FIG. 1: Schematic of a statistical model of antigen recognition by the adaptive immune system. After infection, antigen a encounters immune receptor r at random with a rate $\lambda_a(t)$. An encounter leads to a successful recognition with a probability $f_{r,a}$ that reflects the matching between a given antigen-receptor pair.

inflicted to the organism by the attack of an antigen a is then given by this quantity averaged over the distribution of possible response times: $\bar{F}_a = \int_0^{+\infty} dt F_a(t) H_a(t) = \bar{P}_a \int_0^{+\infty} dm F_a[t_a(m)] e^{-m \bar{P}_a}$, where $t_a(m)$, the inverse function of $m_a(t)$, is the amount of time it takes for m encounters to occur between the immune receptors and pathogen a . The result depends on the cost expressed as a function of the effective time m , $F_a[t_a(m)]$, which we denote $F_a(m)$ to simplify notations.

We will consider several specific choices of the effective cost function in Results. Since not all antigens are equally likely, the overall expected cost is this harm averaged over the antigen distribution:

$$\text{Cost}(\{P_r\}) = \langle F \rangle = \sum_a Q_a \bar{F}_a. \quad (1)$$

The need to defend against many antigens at the same time with a limited number of receptors introduces a trade-off. If more receptors recognize an antigen, there are less to protect against other threats.

Our aim here is to propose a general framework for thinking about the repertoire. Thus, we do not explicitly model intracellular communication, cell differentiation, activation of co-factors, coordination of different cell types, avoidance of self-antigens through thymic selection, and the full complexity of the recognition process. The idea is that $F_a(m)$ implicitly summarizes all of these factors in terms of an effective cost. Of course, more detailed modeling of the cost will be possible as we refine our knowledge of the recognition process.

In general the cost function $F_a(m)$ depends on the antigen a , reflecting the various virulences of different pathogens. To simplify, we can assume that the cost function takes the factorized form: $F_a(m) = \mu_a F(m)$, where μ_a is the pathogen-dependent virulence factor, and $F(m)$ describes how all threats develop with time. The cost will then take the form: $\sum_a \mu_a Q_a \bar{P}_a \int_0^{+\infty} dm F(m) e^{-m \bar{P}_a}$. In this expression, the virulence factor μ_a of a pathogen plays the same role as its likelihood Q_a . Some pathogens are rare but very virulent (like anthrax), while others may be common but not very virulent (like the common cold), and an ideal immune system should be able to cope

with both. In our model the overall “dangerousness” of a pathogen is expressed as the product of the two, $\mu_a Q_a$. Therefore, for all practical purposes μ_a can be absorbed into the definition of Q_a , and will be omitted in the rest of the paper.

Given such a model of the recognition process, there exists an optimal adaptive immune system, characterized by the choice of the receptor distribution P_r , that minimizes the expected cost in a given antigenic environment Q_a . The optimal repertoire is found by minimizing the expected cost in Eq. 1 with respect to P_r , subject to constraints of non-negativity ($P_r \geq 0$) and normalization ($\sum_r P_r = 1$). Simple local extremality conditions are sufficient for optimality because our problem can be shown to be convex (see App. B). The condition $\sum_r P_r = 1$ is a normalized version of the constraint that the total number of receptors is limited.

II. RESULTS

A. The optimal repertoire is more uniform than the pathogen distribution

We can now ask how best to distribute the receptors to minimize the cost (Eq. 1) for a given antigenic environment. To begin, we neglect cross-reactivity (later we will see that this is equivalent to looking at the structure of the repertoire at scales larger than the cross-reactivity). In this case antigens and receptors can be associated one by one by a cross-reactivity function $f_{r,a} = 1$ if $r = a$ and 0 otherwise. In this case we can analytically determine the optimal distribution (App. D 2):

$$P_r^* = \max \left[\bar{F}'^{(-1)} \left(-\frac{\lambda}{Q_r} \right), 0 \right], \quad (2)$$

where $\bar{F}'^{(-1)}$ denotes the inverse function of the derivative of $\bar{F}_a = \bar{F}(\tilde{P}_a)$ expressed as a function of \tilde{P}_a , and λ is a positive constant fixed by the normalization $\sum_r P_r^* = 1$. Table I presents results for several representative cost functions.

A simple scenario occurs when the pathogen population grows exponentially in time, as do the cost and the encounter rate—reflecting the proliferative nature of pathogens. In this case the cost grows linearly in the number of encounters, *i.e.* $F(m) = m$ (see App. C). Then we find that the optimal fraction of the repertoire taken up by a given receptor is proportional to the square root of the frequency of the corresponding antigen $P_r^* \propto \sqrt{Q_r}$. Intuitively, we expect that the optimal repertoire should focus its resources on receptors recognizing the most common antigens. However this enhanced protection against frequent antigens comes at the cost of a slower response against the uncommon antigens, and this bias towards common antigens must remain limited. The square root dependence reflects a particular trade-off between these two opposing constraints, by directing more resources to-

wards common antigens while uniformizing the distribution compared to a linear dependence. Intriguingly, the same square root dependence has been found as an optimal solution for the size of tRNA pools as a function of codon usage [10], and in a model for the screening of suspicious individuals [11].

The extent to which more resources are directed towards common antigens depends on the relative gains and losses of earlier and later recognition events, which are captured in our model by the effective cost function $F(m)$. In general, steeper cost functions imply more flattened distributions of receptors. The cost function $F(m) = m^\alpha$, and its associated optimal distribution $P_r^* \propto Q_r^{1/(1+\alpha)}$, help illustrate this point. Such cost functions can arise when both $m(t)$ and $F(t)$ increase exponentially as a function of time, but with different exponents (see App. C). When α is large, the cost of non recognition increases very quickly with time, calling for an urgent response. Consequently the optimal immune system tends to cover the space uniformly to get all potential threats, even the unlikely ones, under control. Conversely, when α is low, the harm caused by pathogens does not explode with time, meaning that the system can afford to recognize the rarer pathogens late, and focus its resources on the common ones.

In some situations, there may be little or even no difference between a late response, or no response at all, because the total harm caused by an infection stabilizes. For example, consider the cost $F(m) = 1 - e^{-\beta m}$ which saturates at large effective times. In this case, the optimal solution (Table I) relates receptor and antigen through a square root as for linear cost, but with a cut-off at low probabilities. This cut-off occurs because there is little benefit to having receptors recognizing rare antigens, whose recognition is likely to happen late, when differences in recognition times do not matter anymore.

$F(m)$	P_r^*	$A(\tilde{N}_a)$
m^α	$C Q_r^{\frac{1}{1+\alpha}}$	$C' (N_{\text{st}}/\tilde{N}_a)^{1+\alpha}$
$\ln m$	$C Q_r$	$C' (N_{\text{st}}/\tilde{N}_a)$
$1 - \exp(-\beta m)$	$\max\{C\sqrt{Q_r} - \beta, 0\}$	$C' / (\beta + \tilde{N}_a/N_{\text{st}})^2$
$\Theta(m - m_0)$	$\max\{\ln(Q_r)/m_0 - C, 0\}$	$C' \exp(-m_0 \tilde{N}_a/N_{\text{st}})$

TABLE I: The cost function $F(m)$ measures the harm caused to an organism by the time that immune receptors have had m encounters with a pathogen. The optimal receptor distribution P^* is determined by minimizing this cost, given a pathogen distribution Q , and a cross-reactivity function $f_{r,a}$ specifying the probability that receptor r binds to antigen a . The second column gives the form of P^* over scales larger than the cross-reactivity. The optimal P^* can be reached as a steady-state resulting from competitive binding between receptors and antigens (see last section of Results) quantified by an “availability function” A . $\tilde{N}_a = \sum_r N_r f_{r,a}$ represents the coverage of antigen a by the repertoire, $N_{\text{st}} = \sum_r N_r$ is the total steady state population and C, C', β , and m_0 are positive constants.

Real harm may occur only when the effective time m

crosses a threshold. This situation can be modeled by taking $F(m) = \Theta(m - m_0) = 0$ for $m < m_0$, and 1 otherwise. In this case the receptor distribution should be organized to maximize the chance of detection before m_0 . The optimal repertoire for this cost (Table I) has no receptors for the least frequent pathogens (cutoff at low probabilities) and a drastically flattened receptor distribution (logarithm of the pathogen distribution).

Is there a cost function for which the receptor distribution is not flattened relative to the pathogen distribution? This occurs in a special case where cost increases very slowly (logarithmically) with effective time. However, in general, cost is minimized by a receptor distribution that is flattened relative to the pathogen distribution.

B. Cross-reactivity dramatically reduces diversity in the optimal repertoire

By allowing receptors to bind to a variety of antigens, cross-reactivity should permit the immune system to reduce the number of receptor types required to cover the whole range of possible threats. We will show that given sufficient cross-reactivity, the optimal immune repertoire concentrates all its resources on a few receptors, which together tile antigenic space.

Following Perelson and Oster [12], we think of receptors and antigens as points in a common high dimensional *shape space*, whose coordinates are associated to unspecified physicochemical properties. For simplicity, assume that cross-reactivity only depends on the relative position of receptor and antigen in shape space $f_{r,a} = f(r - a)$, where f is a decreasing function of the distance between a and r . Short distances in shape space correspond to a good fit between the two molecules, leading to strong recognition, while large distances translate into weak interactions and poor recognition.

In order to build intuition, we first consider an analytically solvable example (Fig. 2). We describe the space of receptors and antigens by a single continuous number, and assume a Gaussian antigen distribution with variance σ_Q^2 , and Gaussian cross-reactivity of width σ , which sets the typical distance within which a receptor and antigen interact. We derive the optimal receptor distributions analytically for costs of the form $F(m) = m^\alpha$ (App. D 3 b). For narrow cross-reactivities ($\sigma < \sigma_c = \sigma_Q \sqrt{1 + \alpha}$), the optimal receptor distribution is Gaussian with variance $(1 + \alpha)\sigma_Q^2 - \sigma^2$ and the optimal cost is independent of σ . For wide cross-reactivities ($\sigma > \sigma_c$), the receptors are optimally of a single type with reactivity centered on the pathogen distribution, while the optimal normalized cost increases with σ since the receptor is unnecessarily broadly reactive. These results arise from a tension between two opposing tendencies. As in the non cross-reactive case, the need to cover rare pathogens broadens the optimal receptor distribution relative to the pathogen distribution. But cross-reactivity has the opposite effect, favoring more concentrated distributions.




FIG. 2: The optimal cost and receptor distributions for protecting against a one-dimensional Gaussian antigenic landscape $Q(a)$ of variance σ_Q^2 , as a function of the cross-reactivity width σ . As σ increases, the optimal distribution $P^*(r)$ becomes narrower and narrower (left and middle insets), until it concentrates entirely onto a single point, for $\sigma \geq \sqrt{2}\sigma_Q$ (right inset). The minimal cost (multiplied by σ for a comparison at constant recognition capability) is constant below the transition point, but increases with σ past it. The cross-reactivity function, which quantifies the affinity between receptor r and antigen a as a function of their distance in shape space, has a Gaussian form: $f(r - a) = \exp[-(r - a)^2/2\sigma^2]$, and the cost function is linear in the effective recognition time, $F(m) = m$.

Does cross-reactivity generically drive the optimal receptor distribution to cluster into peaks? We investigated this question numerically. For concreteness, we consider a linear cost $F(m) = m$, and random pathogen environments in one or two dimensions constructed by drawing each Q_a from a log-normal distribution characterized by a coefficient of variation κ . For numerical purposes, the shape space is taken to be bounded and discretized, and we use accelerated gradient projection optimization (App. E). We find that the optimal repertoire P^* is strongly peaked on a discrete forest of receptors (Fig. 3A,B). The width of these peaks decreases as numerical precision is increased, suggesting that the true optimum consists of a weighted sum of Dirac delta functions, *i.e.* distinct, discretely spaced receptors in different amounts (see Fig. 6). By inspection, the peaks are spaced evenly, at roughly the cross-reactivity scale σ , suggesting that P^* is smooth when viewed at scales larger than σ . Confirming this, \tilde{P}^* (*i.e.* the coverage of the antigenic space by the receptors) smoothly tracks the variations in the antigen distribution Q_a at a broad scale (Fig. 3A). When viewed coarsely in this way, cross-reactivity is irrelevant and P^* tends to the solutions of Table I.

How can we quantitatively understand such distributions which are fragmentary in detail, but show organization when viewed coarsely (Fig. 3B)? A useful technique, borrowed from condensed matter physics, is to measure the *radial distribution function* [13]: $g(R) = \langle P(r)P(r') \rangle_{|r-r'|=R}$, where $|r - r'|$ is the distance between points r and r' . Fig. 3C presents $g(R)$ for P^* in two dimensions. The initial drop at small r indicates that peaks in P^* are rarely close – *i.e.*, peaks in the optimal

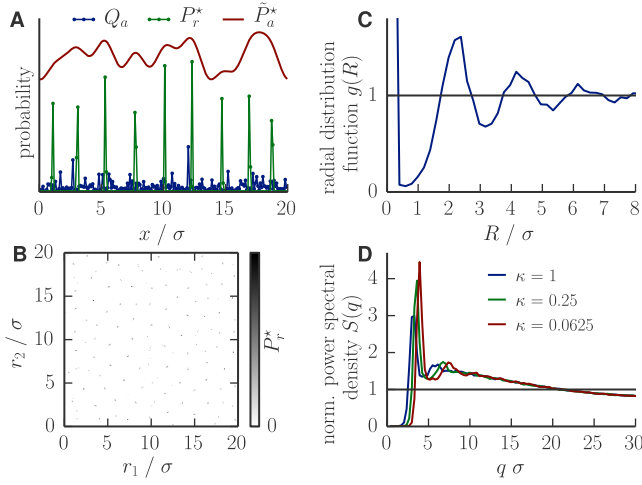


FIG. 3: Cross-reactivity plays an important role in shaping the optimal repertoire, often leading to highly peaked repertoires. (A)-(B): The optimal receptor distribution P_r^* for (A) one- and (B) two-dimensional random environments. Despite being peaked, the optimal distribution of receptors covers the antigenic space fairly uniformly, as shown by its coverage by the receptors, $\tilde{P}_a^* = \sum_r f_{r,a} P_r^*$, shown in the one-dimensional case (A). The cross-reactivity and cost functions are the same as in Fig. 2. The antigenic landscape Q_a is generated randomly from a log-normal distribution with coefficient of variation $\kappa = 1$. (C)-(D): Structural analysis of the tiling pattern formed by the peaks of the optimal receptor distribution P_r^* , in two dimensions. (C) The radial distribution function of P_r^* shows an exclusion zone around each peak, followed by oscillations characteristic of a local tiling pattern. (D) Normalized power spectral density $S(q)$ of P_r^* for different values of the parameter κ quantifying the heterogeneity of the antigenic landscape. The high suppression of fluctuations at large scales (small q) indicates that the pattern has very little fluctuations in the number of receptors used to cover large surface areas.

repertoire tend to repel each other. This exclusion, which operates over the range of strong cross-reactivity, is a sensible way to distribute resources, as it limits redundant protection against the same pathogens. The damped oscillation of the peaks of $g(R)$ confirm that the receptors in P_r^* are organized into a disordered tiling pattern. A similar radial distribution function is seen in high density random packings of hard spheres where the spheres must cover as much space as possible but exclude each other. In both cases, the tiling ensures uniform coverage of space at large scales.

To quantify the regularity of the tiling, we calculate the *normalized power spectral density* of the 2D pattern: $S(q) = \sum_{r,r'} P_r P_{r'} e^{iq(r-r')} / \sum_r P_r^2$, where q is a wave vector. Large (small) $|q|$ correspond to short (long) distances in antigen shape space. When P_r is made of Dirac delta peaks of uniform heights, $S(q)$ coincides with the *structure factor* familiar in physics, and satisfies $S(q \rightarrow \infty) = 1$. Fig. 3D shows $S(q)$ averaged over




FIG. 4: Two individuals in the same environment Q_a that see it with slightly different noises have similar coverages of the antigenic space, but achieve it with different receptors. This results in largely non-overlapping repertoires. Shown are the overlaps (normalized to be between 0 and 1) between the experienced pathogen distributions Q_a , the resulting optimal receptor distributions P_r^* , and the corresponding coverages \tilde{P}_a^* , as a function of the noise ϵ with which individuals perceive the environment. The right plots show an example of antigenic environments and optimal receptor distributions for $\epsilon = 0.2$. We calculated the optimal receptor distributions for two individuals 1 and 2 experiencing respective environments Qe^{z_1} and Qe^{z_2} , where Q is a random environment with fluctuations on scales larger than the cross-reactivity σ (power spectrum $\propto 1/(1 + (10q\sigma)^2)$) normalized so that its coefficient of variation is 0.5, and z_1, z_2 are Gaussian noises of mean zero and variance ϵ^2 . The choice of cost and cross-reactivity functions are the same as in Fig. 2.

many realizations of the antigen landscape, and over all directions of q so that it only depends on its modulus $|q|$. $S(q)$ approaches 1 for large q , showing that the precise local positions of the peaks are random. (The small departure from 1 is attributable to numerical discretization.) $S(q)$ is very low for small q , indicating that the number of receptors contained in any given large area of the shape space is very reproducible, providing uniform coverage. This phenomenon of small scale randomness with large-scale regularity is called *hyperuniformity* [14], and arises in jammed packings [15, 16] as evidence of the incompressibility of the material. For our optimal repertoires small scale fluctuations (large q) get smoothed out by cross-reactivity and can be tolerated, while at large scales the fluctuations track variations in the antigenic landscape to provide smooth coverage (see Fig. 7).

To test the generality of our findings we tested other choices of cross-reactivity functions. We found that the optimal repertoire remains strongly peaked, although the position, number and strength of the peaks do change (Fig. 8). Next we considered distributions of antigens with correlations across shape space (reflecting *e.g.* phylogenetic correlations between pathogens). Again we find peaked optimal receptor distributions (Fig. 9), similar to those for uncorrelated antigen landscapes. For computa-

tional reasons, we restricted our analysis to two dimensional pathogen landscapes, but the analogy with random packing problems that we discussed above allows us to expect that all of these results will hold generally in higher dimensions.

In summary, the optimal immune repertoire looks random at scales smaller than the cross-reactivity, but has the structure of a disordered tiling at larger scales so that, after accounting for cross-reactivity, the repertoire smoothly covers the pathogen landscape. These findings have an important consequence for different individuals exposed to the same pathogenic environment. Each individual will experience a slightly different spectrum of antigens because of the statistics of encounters and other sources of variability. These slightly different experiences of the same world lead to optimal repertoires with a striking property – the receptor distributions are largely different, even though their coverage of the pathogen landscape is similar after including cross-reactivity (Fig. 4). This finding can be compared with surveys of “public” repertoires of immune receptors [2, 17].

C. The optimal repertoire can be reached through competition for antigens

The results presented so far have established how repertoires should be structured to provide optimal protection. Given the complex interdependences between receptors arising from local and global trade-offs, one might think that the globally optimal solution could only be reached via some biologically implausible centralized mechanism distributing resources system-wide. In fact, we will show that the optimal repertoire can be reached through self-organization, via competitive evolution of receptor populations under antigen stimulation.

We consider a model that is similar to that introduced by de Boer, Perelson and collaborators for competitive dynamics of B and T cells [18, 19]. Its main assumptions are that division of receptor-expressing lymphocytes is driven by antigen stimulation, and that receptors compete for the limited supply of antigens. The number N_r of receptors of a given type r evolves according to:

$$\frac{dN_r}{dt} = N_r \left[\sum_a Q_a A \left(\sum_r N_r f_{r,a} \right) f_{r,a} - d \right]. \quad (3)$$

Receptors proliferate upon successful recognition by antigens (first term of the equation) and die with a constant rate d (second term of the equation). The growth rate of a receptor type is proportional to the number of antigens that it detects. In the absence of competition, this amount is simply $\sum_a Q_a f_{r,a}$, but the antigen a will also bind other receptors, reducing its availability for receptor r . The *coverage* of antigen a by the repertoire, $\tilde{N}_a = \sum_r N_r f_{r,a}$, quantifies the breadth of the receptor pool competing to bind with a . The availability of antigen a for binding is assumed to be a decreasing function

$A(\tilde{N}_a)$ of its coverage. The stimulation of r by a is thus modified to: $\sum_a Q_a A(\tilde{N}_a) f_{r,a}$ as in Eq. 3. For a given pathogenic environment, the total steady-state receptor population size N will be set by the death rate d , which counter-balances growth at steady state.

The stable fixed points of the dynamics (3) realize the optimal repertoires of the previous sections when the availability function A is matched to the cost function $F(m)$ through the relation

$$A(\tilde{N}_a) = -c' \bar{F}'(\tilde{N}_a/N_{\text{st}}), \quad (4)$$

where N_{st} is the total number of receptors $\sum_r N_r$ at steady state. Table I shows $A(\tilde{N})$ for several cost functions. To understand this result, first note that when binding is *not* cross-reactive the dynamical equations for each receptor are independent, and read: $dN_r/dt = N_r(Q_r A(N_r) - d)$. The availability function now depends only on N_r , meaning that receptors only compete with their own kind — they occupy their own antigenic niche. The steady state size of clone r is thus set by the carrying capacity of that niche, $N_r = A^{(-1)}(d/Q_r)$, or zero if that capacity is negative. With the availability given by Eq. 4, this reproduces the optimal repertoire (Eq. 2). A similar argument holds when receptor binding is cross-reactive (App. F). Cross-reactivity leads to competition amongst receptor types, effectively enforcing an exclusion between similar receptors. This phenomenon, known in ecology as competitive exclusion, is important for lymphocyte dynamics [18], and provides the mechanism by which our dynamical model reproduces the discrete clustering found in the optimal receptor distribution.

To check that the dynamics do converge to the optimum, we simulated Eq. 3 numerically for a random antigenic environment in two dimensions, with $A(\tilde{N}) = 1/(1 + \tilde{N}/N_0)^2$. Fig. 5 shows the dynamics of the receptor distribution $P_r(t) = N_r(t)/\sum_{r'} N_{r'}(t)$, as well as its cost relative to the optimal solution, as a function of time. Starting from a uniform initial distribution of receptors, the repertoire reorganizes into localized peaks that become increasingly prominent and well-separated with time, with almost no receptors in between. Starting from a random initial condition leads to the same steady state (Fig. 10). The cost converges towards the global minimum, indicating that the steady-state solution is indeed optimal.

In summary, competitive dynamics can allow the immune repertoire to self-organize into a state that confers high protection against infections. In the special case when the availability A is scale invariant, the expected cost is a Lyapunov function of the dynamics (App. G). In this case, we can prove analytically that regardless of the initial condition the cost will steadily decrease until the optimum is reached.

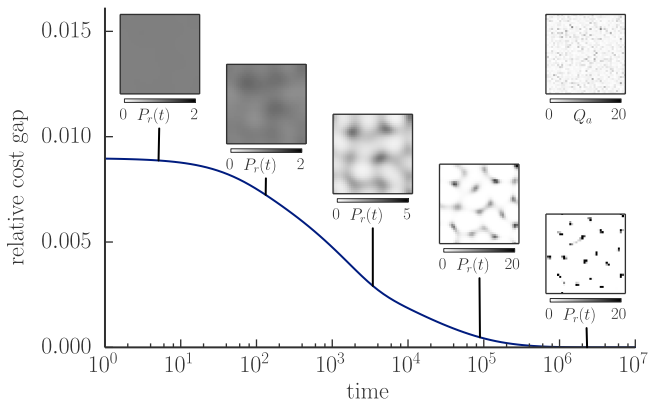


FIG. 5: The immune repertoire can self-organize to a state that minimizes cost and provides protection against infections via competitive evolution of receptor populations stimulated by antigens. Numerical solution of the population dynamics (Eq. 3) shows how competition causes a uniform initial receptor distribution to fragment into a highly peaked pattern (insets representing $P_r(t) = N_r(t)/\sum_{r'} N_{r'}(t)$). The top-right inset represents the antigenic environment Q_a driving the dynamics (generated as in Fig. 3B). Departure from optimality, as measured by the relative cost gap $[\langle F \rangle(P_r(t)) - \langle F \rangle(P_r^*)]/\langle F \rangle(P_r^*)$, decreases with time and eventually reaches zero. We use the availability function $A(\tilde{N}) = 1/(1 + \tilde{N}/N_0)^2$ with $N_0 = 10^6$, a death rate $d = 0.01$ and a cost function $F(m) = 1 - e^{-\beta m}$ with $\beta = 0.04$. The space size is 10σ . The initial condition is uniform with $\sum_r N_r(0) = 2.5 \cdot 10^7$.

III. DISCUSSION

We introduced a general framework for predicting the optimal composition of the immune repertoire to minimize the cost of infections contracted from a given distribution of antigens. This framework can be extended in several ways to be more biologically faithful, *e.g.* by accounting for receptor-dependent cross-reactivities, antigen-dependent infection dynamics, and evolution of the pathogenic landscape. Our predictions can be tested in experiments that study how the environment influences the composition of immune repertoires, either via high-throughput sequencing surveys of receptor populations [2, 20], or by sequencing receptors specific to given antigens [21]. The comparison between theory and experiment will provide insight into the functional constraints of antigen recognition by the immune system.

There are many situations where living systems must respond to very diverse and often very high dimensional spaces of external influences using strictly limited resources. To sense, internally represent, and then respond to these influences, organisms often employ a large diversity of components, such as cell types or genes [22], each sensitive to a small part of the space. For example, the retina supports a diverse population of ganglion cell types, each sensitive to a different visual feature, that

collectively represent the behaviorally salient aspects of visual scenes [23, 24]. Likewise, the mammalian olfactory system contains some ~ 1000 distinct receptors that each bind widely to odorants, and collectively cover olfactory space [25]. In these cases, the limited repertoire of component types provides a key constraint on information processing. Faced with such constraints, living systems must commit resources wisely, adapting to the structure of the environment, and balancing breadth of coverage against depth of resolution, in light of priorities, costs and constraints [26]. We have shown that these elements also shape the optimal form of the immune repertoire.

Our finding that cross-reactivity causes the optimal repertoire to fragment is related to the concept of *limiting similarity* due to competitive exclusion in ecological settings [27–31]. In the latter context, empty regions of phenotypic space result when competition is important on the scale at which resources vary [27], and continuous coexistence of species only occurs in exceptional cases [30]. In general, niche-space heterogeneity promotes species clustering [28, 31], recalling our finding that any heterogeneous antigen distribution leads to fragmentation of the optimal repertoire. The conceptual connection between the immune repertoire and ecological organization is even clearer in our dynamical model where species compete for an array of resources (the antigens), and grow in relation to their success in securing resources.

Although this study relies on a simple abstraction of the adaptive immune system, we expect that our framework and results will extend to other distributed protection systems where diverse threats are addressed by an array of specific responses. For example, the immune system of bacteria, or CRISPR system [32], for which population dynamics models have already been proposed [33], could be studied within a similar framework to predict the relative abundance of CRISPR spacers and corresponding viruses in a co-evolving population of bacteria and viruses.

Acknowledgements. The work was supported by grant ERCStG n. 306312. VB was supported by the Fondation Pierre-Gilles de Gennes, NSF grants PHY-1058202 and EF-0928048. Portions of this work were done at the Aspen Center for Physics, supported by NSF grant PHY-1066293. AM was supported by a DAAD Promos stipend.

Appendix A: Probability distribution of the time of first recognition

In order to calculate the cost of not-recognizing an antigen a , we need to find the distribution of times when a successful encounter takes place. The probability of having the first recognition of antigen a by receptor r in the

time between t and $t + dt$ reads:

$$H_a(t)dt = \lambda_a(t)dt \cdot \sum_r P_r f_{r,a} \\ \times \lim_{N \rightarrow \infty} \prod_{i=1}^N \left(1 - \lambda_a(t_i) \frac{t}{N} \sum_r P_r f_{r,a} \right),$$

where the first term is the probability of having an encounter between t and $t+dt$, the second the probability of this encounter being successful, and the third the probability of there not being any prior recognition events. For the calculation of the last term we have decomposed the time leading up to t into N intervals of length t/N . Taking the $N \rightarrow \infty$ limit yields:

$$H_a(t) = \lambda_a(t) \tilde{P}_a e^{-\int_0^t dt' \lambda_a(t') \tilde{P}_a}, \quad (\text{A1})$$

where we have used the short-hand notation $\tilde{P}_a = \sum_r P_r f_{r,a}$ for the probability that a randomly chosen receptor recognizes antigen a .

Appendix B: Convexity of the expected cost

In this Appendix we show that the cost function $\langle F \rangle$ is a convex function of its argument $\{P_r\}$ (the receptor distribution). We start by introducing an alternative expression of \bar{F}_a , obtained by integration by parts:

$$\bar{F}_a = \int_0^\infty dm F'_a(m) e^{-m \tilde{P}_a} + F(0). \quad (\text{B1})$$

We calculate the derivatives of this average cost with respect to \tilde{P}_a :

$$\frac{d\bar{F}_a}{d\tilde{P}_a} = - \int_0^\infty dm m F'_a(m) e^{-m \tilde{P}_a} \quad (\text{B2})$$

$$\frac{d^2 \bar{F}_a}{d\tilde{P}_a^2} = \int_0^\infty dm m^2 F'_a(m) e^{-m \tilde{P}_a} \quad (\text{B3})$$

Since by assumption $F'_a(m)$ is positive, the second derivative of \bar{F}_a with respect to \tilde{P}_a is positive. This establishes the convexity of \tilde{P}_a as a function of \bar{P}_a . Since $\langle F \rangle = \sum_a Q_a \bar{F}_a$ (with $Q_a \geq 0$), it is a convex function of $\{\tilde{P}_a\}$. Therefore it is also a convex function of $\{P_r\}$, as $\{P_r\}$ and $\{\tilde{P}_a\}$ are linearly related.

Appendix C: Biological motivation of power-law cost functions

In the main text we have developed a general framework for discussing the antigen-receptor recognition process. To fully specify the model we need to choose an effective cost function $F_a(m) = F_a(t_a(m))$. In the main text we derive optimal receptor distributions for a number of effective cost functions, including power-law functions $F(m) = m^\alpha$. Here we sketch plausible scenarios motivating that choice.




FIG. 6: Solving the optimization problem with a finer and finer discretization step suggests that the peaks found in the optimal receptor distributions converge to true Dirac delta functions. Starting from a problem with a discretization step of $\Delta = 0.1\sigma$, we construct coarse-grained versions of it by downsampling the antigen distribution two and four fold, yielding $\Delta = 0.2$ and 0.4 respectively. The resulting coarse-grained optimization problems are then solved, and the optimal distributions P_r^*/Δ represented (after appropriate normalization by the step size). The random antigen distribution is log-normal with coefficient of variation $\kappa = 0.25$.

Consider an organism being infected with a antigen a . As long as there is no immune reaction, the antigens divide inside its host and thus increase its population size. If the initial population size is small it is reasonable to assume exponential growth.

The more antigens there are at the time of the immune reaction the more damage they can potentially do. Likewise, the more antigens, the higher the rate of encounters. These two quantities are also expected to grow exponentially in time:

$$F_a(t) = F_a(0) e^{\nu_a t}, \quad (\text{C1})$$

$$\lambda_a(t) = \lambda_a(0) e^{\nu'_a t} \quad (\text{C2})$$

The two exponents may be different in general, because the number of pathogenic agents that cause the harm may grow differently than the number of antigens that can be recognized by the immune system. This difference could for example come stem from the fact that both the pathogen's antigenic exposure and its virulence are cooperative effects, and thus scale as a power of the number of invading individuals. Using $m_a(t) = \lambda_a(0)(e^{\nu'_a t} - 1)/\nu'_a$, and eliminating time $t \approx \ln[m_a/\lambda_a(0)]/\nu'_a$ (for t large compared to $1/\nu'_a$), we rewrite the effective cost function in terms of the number of encounters:

$$F_a(m) = F_a(0) \left(\frac{m}{\lambda_a(0)} \right)^{\frac{\nu_a}{\nu'_a}} \propto m^\alpha, \quad (\text{C3})$$

with $\alpha = \nu_a/\nu'_a$.




FIG. 7: Power spectral density normalized by the squared antigenic environment heterogeneity index κ : $|\sum_r P_r e^{iqr}|^2 / \kappa^2$. The data collapse for different κ shows that the fluctuations at large scale are entirely attributable to those of the antigenic environment, and scale with them. At these large scales, the power spectrum of the receptor distribution is approximately given by: $\exp[(q\sigma)^2/2]/4$. The exponential term stems from the inverse of the Fourier transform of f (see Eq. D21). Parameters are the same as in Fig. 3.

Appendix D: Analytical optimization

1. Optimality conditions

In the following we give optimality conditions for the optimization problem defined in the main text, which will be used for the following analytical determination of optimal receptor distributions. These conditions, called Karush-Kuhn-Tucker conditions [34], are derived from a generalization of the method of Lagrange multipliers to inequality as well as equality constraints.

The Lagrangian for the optimization problem is

$$\mathcal{L}(P, \lambda, \nu) = \langle F \rangle(P) + \lambda \left(\sum_r P_r - 1 \right) - \sum_r \nu_r P_r, \quad (\text{D1})$$

with

$$\langle F \rangle = \sum_a Q_a \bar{F}_a. \quad (\text{D2})$$

λ is a Lagrange multiplier enforcing the normalization constraint and ν_r are Lagrange multipliers enforcing the non-negativity constraint. The optimal P^* is an extremum of this Lagrangian. Therefore the stationarity conditions:

$$\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} + \lambda^* - \nu_r^* = 0, \quad (\text{D3})$$

with

$$\frac{\partial \langle F \rangle}{\partial P_r} = \sum_a Q_a \bar{F}'_a(\bar{P}_a) f_{r,a}, \quad (\text{D4})$$

must hold for some value of λ^* and ν_r^* that enforce the constraints. The inequality constraint $P_r \geq 0$ further requires that:

$$\nu_r^* \geq 0 \quad (\text{D5})$$

$$\nu_r^* P_r^* = 0, \quad (\text{D6})$$

where the second is known as the complementary slackness condition. It requires the Lagrange multipliers associated with the non-negativity to be zero unless the constraint is active, *i.e.* unless the corresponding receptor probability is zero.

The three conditions may be reformulated as:

$$\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} + \lambda^* \geq 0 \quad (\text{D7})$$

$$\left(\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} + \lambda^* \right) P_r = 0 \quad (\text{D8})$$

For all receptors that are present in the optimal repertoire ($P_r^* > 0$) these conditions imply

$$\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} = -\lambda^*. \quad (\text{D9})$$

If a receptor is not present in the optimal repertoire ($P_r^* = 0$) then the less stringent condition holds:

$$\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} \geq -\lambda^*. \quad (\text{D10})$$

We note here that $\partial \langle F \rangle / \partial P_r \leq 0$ (because more receptors always yield a lower cost), so that $\lambda^* \geq 0$.

These two conditions can be explained as follows: if a repertoire is optimal, all changes allowed by the constraints will lead to a higher cost, *i.e.* moving receptors from one type to another will not yield an improvement. All partial derivatives of the cost with respect to the receptor probabilities should thus be equal to the same value (Eq. D9). If there are already no receptors of a certain type, *i.e.* $P_r = 0$, we get a less stringent condition. We can no longer remove receptors away from this type r , but only add some to it, at the expense of other receptor types. The increase in cost due to the depletion of these other types should be higher than the gain of moving them to type r . The partial derivatives of the cost with respect to the receptors that are not present in the repertoire must thus be larger than the partial derivatives of the present receptors, which are given by $-\lambda^*$ (Eq. D10).

2. Solution for uniquely specific receptors

We now solve Eqs. D9 and D10 for a repertoire of uniquely specific receptors (no cross-reactivity). Eq. D4 becomes

$$\frac{\partial \langle F \rangle}{\partial P_r} = Q_r \bar{F}'_r(P_r), \quad (\text{D11})$$

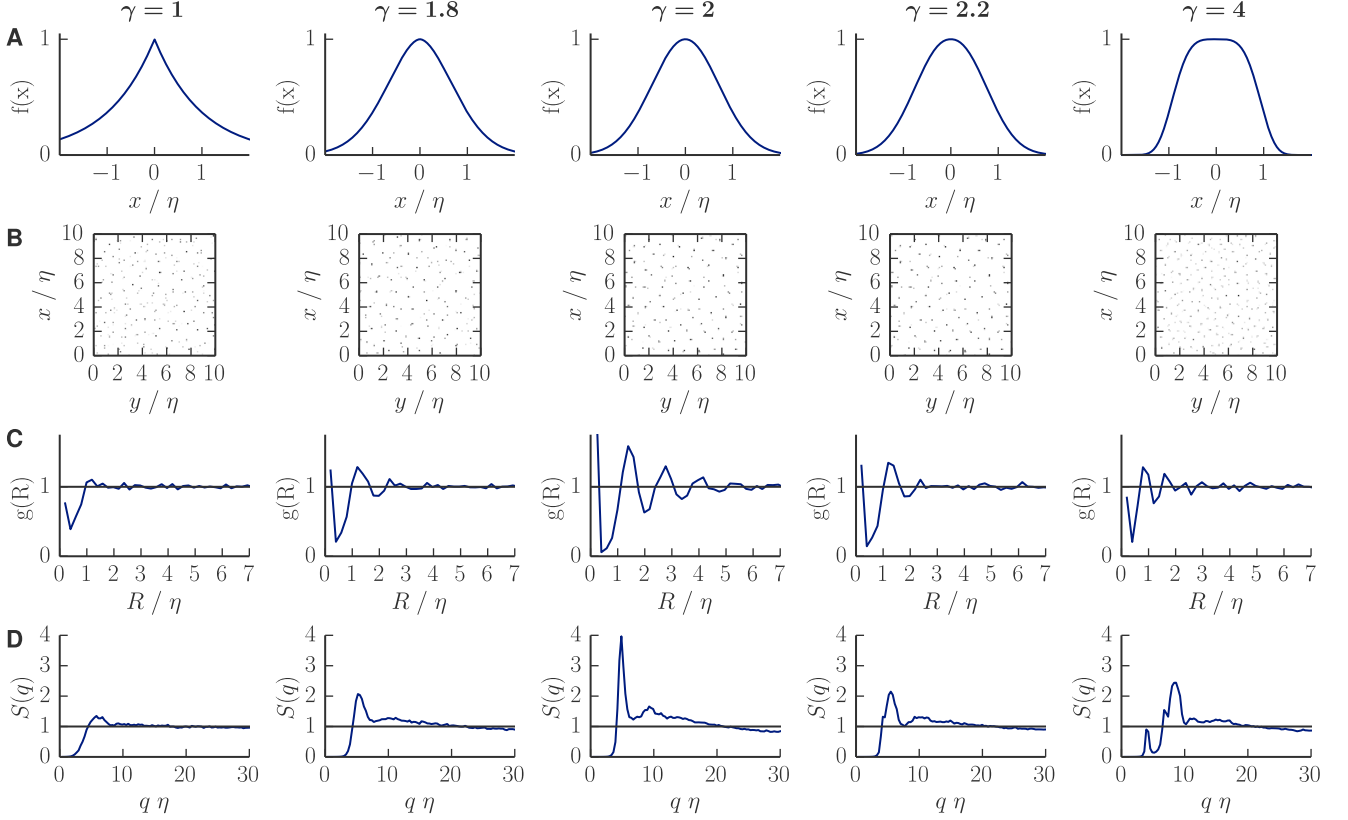


FIG. 8: Influence of the choice of the cross-reactivity kernel $f(a-r)$ on the optimization problem. Regardless of the kernel choice the optimal repertoire is peaked for non-uniform antigen distributions. The details of distribution depend on the cross-reactivity kernel. (A): Kernel functions used to describe cross-reactivity. We use the family of kernel functions defined by $f(r-a) = \exp[-(|r-a|/\eta)^\gamma]$. By changing the parameter γ we can go from an exponential ($\gamma = 1$) via a Gaussian $\gamma = 2$ to a top-hat kernel ($\gamma \rightarrow \infty$). Up to $\gamma = 2$ all such kernels have positive Fourier transforms, whereas for $\gamma > 2$ the Fourier transforms also take negative values [30]. (B): Examples of optimal receptor distributions in two dimensions, for antigenic environments generated as in Fig. 3B (with coefficient of variation $\kappa = 0.25$). (C) Radial distribution function of the optimal distribution. (D) Structure factor of the optimal distribution. The results in both (C) and (D) are averaged over 10 independent runs. A linear effective cost function $F(m) = m$ is assumed throughout. The random antigen distribution is log-normal with coefficient of variation $\kappa = 0.25$.

where we have used the fact that in the absence of cross-reactivity $\bar{P}_a = P_a$. If all optimal receptor probabilities are positive then we can insert this relationship into Eq. D9 to obtain

$$Q_r \bar{F}'_r(P_r^*) = -\lambda^*. \quad (\text{D12})$$

and thus:

$$P_r^* = h_r(-\lambda^*/Q_r), \quad (\text{D13})$$

where $h_r = \bar{F}'_r{}^{(-1)}$ denotes the inverse function of \bar{F}'_r . Since that function \bar{F}'_r is always negative, h_r must take a negative argument.

For some cost functions, solving this equation may yield some negative receptor probabilities. In these cases some of the non-negativity constraints need to be active. Setting $P_r = 0$ when Eq. D13 is negative yields the correct optimal distribution under the non-negativity constraint. We verify that for these r , Eq. D10 is satisfied

by $P_r = 0$, because:

$$Q_r \bar{F}'_r(P_r = 0) \geq Q_r \bar{F}'_r[h_r(-\lambda^*/Q_r)] = -\lambda^*, \quad (\text{D14})$$

where we have used the fact that \bar{F}'_r is a increasing function of its argument (due to the positivity of its derivative, cf. Eq. B3), and $h_r(-\lambda^*/Q_r) \leq 0$

In summary, the solution to the optimization problem is

$$P_r^* = \max\{h_r(-\lambda^*/Q_r), 0\}, \quad (\text{D15})$$

where the value of λ^* is fixed by the normalization condition $\sum_r P_r = 1$.

In Tab. II we give the explicit expressions of \bar{F}_a and h_a , for the particular choices of the cost function $F(m)$ considered in the main text.




FIG. 9: Adding correlations to the antigen distribution does not change the peakedness of optimal receptor distributions. The result of the optimization is shown for a random antigen landscape with correlations. The antigen distribution is generated by Fourier filtering. First we generate an uncorrelated, normally distributed random series. This series is then filtered to obtain a power spectrum $\propto 1/(1 + (10q\sigma)^2)$. Finally, the filtered series is exponentiated to ensure the non-negativity of the generated values.




FIG. 10: Numerical solution of the population dynamics described in the main text as a function of time. The same calculation as in Fig. 5, but using a random initial condition (log-normal with coefficient of variation 1) shows the same convergence to the optimal receptor distribution.

$F(m)$	$\bar{F}(\tilde{P}_a)$	$h(x)$
m^α	$\Gamma(1 + \alpha)/\tilde{P}_a^\alpha$	$(-x/(\alpha\Gamma(1 + \alpha)))^{1/(1+\alpha)}$
$\ln m$	$\gamma - \ln \tilde{P}_a$	$-1/x$
$1 - \exp(-\beta m)$	$\beta/(\beta + \tilde{P}_a)$	$\sqrt{-\beta/x} - \beta$
$\Theta(m - m_0)$	$\exp(-m_0\tilde{P}_a)$	$-\ln(-x/m_0)/m_0$

TABLE II: Intermediate results in the derivation of the optimal solution. The first column shows several choices of the effective cost function, $F(m)$. For these cost functions the second column shows the average cost of a pathogenic attack, $\bar{F}(\tilde{P}_a)$, and the third column shows the inverse of its derivative, $h = (\bar{F}')^{-1}$. Γ is the Gamma function, γ is Euler's constant, β and m_0 are positive constants.

3. Solution for cross-reactive receptors

The previous results can be generalized to cross-reactive receptors in a continuous space, using Fourier transforms. This generalization will lead up to the results presented in the *Cross-reactivity dramatically limits optimal repertoire diversity* section of the main text, and notably the Gaussian case discussed therein.

a. Deconvoluting the optimality conditions in Fourier space

We consider a continuous receptor-antigen space and we assume a translation invariant cross-reactivity function $f_{r,a} = f(r - a)$. We write the optimality condition Eq. D9

$$\int dp Q(a) \bar{F}'[\tilde{P}^*(a)] f(r - a) = -\lambda^*, \quad (\text{D16})$$

where in continuous space the coverage is defined as:

$$\tilde{P}(a) = \int dr P(r) f(r - a). \quad (\text{D17})$$

We notice that both expressions involve integrals, which are convolutions with the cross-reactivity kernel. Since the convolution of a constant is also a constant, a solution of

$$Q(a) \bar{F}'(\tilde{P}^*(a)) = -\lambda', \quad \text{with } \lambda' > 0, \quad (\text{D18})$$

is also a solution of Eq. D16. As in the case of uniquely specific receptors, we can solve this equation for $\tilde{P}^*(a)$:

$$\tilde{P}^*(a) = h[-\lambda'/Q(a)], \quad (\text{D19})$$

where $h = \bar{F}'^{(-1)}$ as in D13. If there was no cross-reactivity, there would be no difference between P and \tilde{P} , and we would be done. Here we need to perform a deconvolution to obtain the optimal receptor distribution P from the optimal coverage \tilde{P} . We do so in Fourier space, where the convolution turns into a product. Deconvolution is therefore much simpler in Fourier space as it corresponds to a division

$$\mathcal{F}[\tilde{P}] = \mathcal{F}[P]\mathcal{F}[f] \Leftrightarrow \mathcal{F}[P] = \mathcal{F}[\tilde{P}]/\mathcal{F}[f], \quad (\text{D20})$$

where we have defined the Fourier transform of a function $g(x)$ as $\mathcal{F}[g](k) = \int_{-\infty}^{\infty} dx g(x) e^{ikx}$. To calculate the optimal receptor distribution we insert Eq. D19 into Eq. D20 and perform an inverse Fourier transform $\mathcal{F}^{-1}[\tilde{g}](x) = (1/2\pi) \int_{-\infty}^{\infty} dk \tilde{g}(k) e^{-ikx}$ to obtain

$$P^* = \mathcal{F}^{-1}[\mathcal{F}[h(-\lambda'/Q)]/\mathcal{F}[f]]. \quad (\text{D21})$$

This result is only valid as long as the above quantity is positive and normalizable, as we shall see below.

b. The Gaussian case

In this section we apply the general results of the previous section to a concrete example. In order to find the optimal receptor distribution analytically we use Eq. D21, we assume the antigen distribution and cross-reactivity function are Gaussian

$$Q(a) = \frac{1}{\sqrt{2\pi\sigma_Q^2}} \exp(-a^2/2\sigma_Q^2), \quad (\text{D22})$$

$$f(r-a) = \exp[-(r-a)^2/2\sigma^2], \quad (\text{D23})$$

and we take

$$F(m) = m^\alpha. \quad (\text{D24})$$

Inserting h from Tab. II into Eq. D21 allows us to write

$$P^* \propto \mathcal{F}^{-1} \left[\mathcal{F}[Q^{\frac{1}{1+\alpha}}] / \mathcal{F}[f] \right] \quad (\text{D25})$$

as an equivalent equation determining the optimal repertoire. We can calculate the modified antigen distribution as

$$Q(a)^{\frac{1}{1+\alpha}} \propto \exp\left(-\frac{a^2}{2(1+\alpha)\sigma_Q^2}\right). \quad (\text{D26})$$

The Fourier transform of a Gaussian function of variance σ^2 is a Gaussian function of variance $1/\sigma^2$ [35]. Therefore we have

$$\mathcal{F}[Q^{\frac{1}{1+\alpha}}](q) \propto \exp[-(1+\alpha)\sigma_Q^2 q^2/2], \quad (\text{D27})$$

$$\mathcal{F}[f](q) \propto \exp[-\sigma^2 a^2/2], \quad (\text{D28})$$

from which

$$\mathcal{F}[Q^{\frac{1}{1+\alpha}}] / \mathcal{F}[f] \propto \exp\{-(1+\alpha)\sigma_Q^2 - \sigma^2\} q^2/2\} \quad (\text{D29})$$

follows. Taking the inverse Fourier transform and normalizing, we obtain

$$P^*(r) = \frac{1}{\sqrt{2\pi[(1+\alpha)\sigma_Q^2 - \sigma^2]}} \exp\left(-\frac{r^2}{2[(1+\alpha)\sigma_Q^2 - \sigma^2]}\right). \quad (\text{D30})$$

Normalization is only possible for $\sigma < \sigma_Q \sqrt{1+\alpha} \equiv \sigma_c$. In the limit $\sigma \rightarrow \sigma_c$ the Gaussian converges to a Dirac delta function. Intuition suggests that a Dirac delta function centered on the peak position should remain optimal for further increases in σ . To prove this assertion we note that a Dirac delta function is zero everywhere, except in one point. Since all but one receptor probabilities are at the boundary defined by the non-negativity constraints, we only need to check Eq. D10. We compute the left-hand side of Eq. D16 as a function of r

$$\begin{aligned} & \int dp Q(a) \bar{F}'[\tilde{P}^*(a)] f(r-a) \\ & \propto -\exp\left\{\frac{-r^2[\sigma^2 - (1+\alpha)\sigma_Q^2]}{2\sigma^2(\sigma^2 - \alpha\sigma_Q^2)}\right\}, \end{aligned} \quad (\text{D31})$$

and note that it has a minimum for $r = 0$. This shows that the partial derivatives of the expected cost at $r \neq 0$ are greater than at $r = 0$, implying that Eq. D10 holds.

The cost of the optimal repertoires as a function of the cross-reactivity width σ is given by

$$\langle F \rangle (P^*) = \left(\frac{\sigma_Q}{\sigma}\right)^\alpha \begin{cases} (1+\alpha)^{\frac{1+\alpha}{2}} & \text{if } \sigma < \sigma_c, \\ \frac{(\sigma/\sigma_Q)^\alpha}{\sqrt{1-\alpha(\sigma_Q/\sigma)^2}} & \text{otherwise.} \end{cases} \quad (\text{D32})$$

Both expressions give the same cost at the transition $\sigma = \sigma_c$. After multiplying by $(\sigma/\sigma_Q)^\alpha$ to compare at constant recognition capability $\int f = \sqrt{2\pi}\sigma$, this expression is constant for $\sigma < \sigma_c$, and grows for $\sigma > \sigma_c$.

c. General argument for peakedness

A simple argument can help understand why cross-reactivity generically leads to peaked optimal solutions. The convolution with a kernel is a smoothing operation, represented by a low-pass filter in the Fourier domain. The optimal solution in the absence of the non-negativity constraints requires that $\tilde{P}_a = h(Q_a)$. As \tilde{P}_a is the low-passed filtered version of P_r , the high-frequency components of $h(Q_a)$ will be magnified by the deconvolution. These high-frequency wiggles can lead to negative values of $\mathcal{F}^{-1}[h(Q_a)]$, which are not allowed, leading to set many values of $P(r)$ to zero. This effects results in a peaked solution. Because the size of the cross-reactivity kernel is inversely proportional to the cutoff frequency in the Fourier domain, we expect the spacing of the peaks to be related to the size of the cross-reactivity kernel.

Appendix E: Numerical optimization

We numerically minimize the cost function subject to the normalization and non-negativity constraints by using a fast projected gradient algorithm. In the following we provide details on this numerical algorithm. To facilitate notations let us define the function to minimize as $g(x)$, where x is a vector in a Euclidean space, and the convex set C defined by the constraints. In these notations the problem we want to solve can be stated as

$$\min_{x \in C} g(x). \quad (\text{E1})$$

Given an arbitrary starting point $x_0 \in C$ the algorithm performs the following iterative procedure:

$$y^{k+1} = x^k + \omega^k (x^k - x^{k-1}), \quad (\text{E2})$$

$$x^{k+1} = \mathcal{P}(y^{k+1} - s^k \nabla g(y^{k+1})), \quad (\text{E3})$$

where ∇ denotes the gradient. Here \mathcal{P} denotes a projection onto C , ω^k is an extrapolation step size and s^k is the step size taken in the direction of the gradient. The

extrapolation step size has to be chosen carefully to ensure the faster convergence of this method with respect to an ordinary gradient method. Following [36] we use

$$\omega_k = \frac{k}{k+3}. \quad (\text{E4})$$

The step size s is determined by backtracking [37]: we iteratively decrease s by multiplication by $\beta < 1$ until $g(z) \leq g(y^k) + (x - y) \cdot \nabla g(y^k) + \frac{1}{2s}(z - y)^2$, where $x \cdot y$ denotes the inner dot product between x and y , and $z = \mathcal{P}(y^k - s\nabla g(y^k))$. In practice we determine s in this way at the first step of the optimization and then keep it fixed based on this initial estimate.

The projection of a point onto a convex set is defined by the following quadratic programming problem:

$$\mathcal{P}(y) = \operatorname{argmin}_{x \in C} \frac{1}{2}(x - y)^2. \quad (\text{E5})$$

If the convex set is a simplex as is the case for our problem, there fortunately exist efficient algorithms for solving this problem. We use the algorithm described in [38].

To stop the iteration one needs to define a suitable stopping criteria. As the problem is convex we can establish a lower bound for the cost by solving a linear programming problem as follows:

$$g_{lb} = g(x^k) + \min_{x \in C} [(x - x^k) \cdot \nabla g(x^k)] \leq g(x^*). \quad (\text{E6})$$

The linear programming problem $\bar{x}^k = \operatorname{argmin}_{x \in C} \nabla g(x^k)^T (x - x^k)$ is solved explicitly [39] by

$$\bar{x}^k = e_{i^*}, \quad i^* = \operatorname{argmin}_i (\nabla g(x^k))_i, \quad (\text{E7})$$

where e_i denotes the i th unit vector. We can use this lower bound to define a stopping criterion for the numerical optimization

$$\frac{g(x^k) - g_{lb}}{g_{lb}} < \epsilon. \quad (\text{E8})$$

For all reported numerical results we have chosen $\epsilon = 10^{-8}$.

The discretization steps used in the figures are listed below:

Step	Figure
0.25 σ	5, 10
0.1 σ	3, 7, 8, 9
0.05 σ	4

Appendix F: Stable fixed point of population dynamics minimizes corresponding cost function

In this section we show that the stable fixed point $\{N_r^*\}$ of the population dynamics:

$$\frac{dN_r}{dt} = N_r \left[\sum_a Q_a A \left(\sum_r N_r f_{r,a} \right) f_{r,a} - d \right] \quad (\text{F1})$$

gives a probability distribution $P_r = N_r/N_{\text{tot}}$ (with $N_{\text{tot}} = \sum_r N_r$) that minimizes the cost $\langle F \rangle$. For this correspondence to be exact, the availability function of the dynamics and the effective cost function of the optimization must be related by:

$$A(\tilde{N}_a) = -c' \bar{F}'(\tilde{N}_a/N_{\text{st}}), \quad (\text{F2})$$

where $\tilde{N}_a = \sum_r N_r f_{r,a}$, and N_{st} is the total number of receptors N_{tot} at the fixed point.

A fixed point is characterized by $dN_r/dt = 0$. If $N_r > 0$, this translates into

$$\sum_a Q_a A \left(\sum_r N_r f_{r,a} \right) f_{r,a} - d = 0. \quad (\text{F3})$$

Using the correspondence between availability and cost function given by Eq. F2 we rewrite this condition as

$$\sum_a Q_a \bar{F}'(\tilde{P}_a) f_{r,a} = -c' d, \quad (\text{F4})$$

which is equivalent to the optimality condition Eq. D9, with the identification $\lambda^* = c' d$.

For $N_r = 0$ we need to work a bit harder to show that the optimality condition at the boundary Eq. D10 is satisfied. Here the key assumption establishing the minimization of the cost function is the stability of the fixed point. A fixed point is stable if the real parts of the Jacobian's eigenvalues are all negative. The Jacobian reads:

$$J_{r,r'} = \delta_{r,r'} \left(\sum_a Q_a A \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} - d \right) + N_r \sum_a Q_a A' \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} f_{r',a}. \quad (\text{F5})$$

We remark that for $N_r = 0$ the r^{th} row of the Jacobian is non-zero only on the diagonal. That value on the diagonal is an eigenvalue of the Jacobian and must be negative:

$$\sum_a Q_a A \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} - d < 0, \quad (\text{F6})$$

Again we replace $A(\sum_r N_r f_{r,a})$ by $-\bar{F}'_a(\tilde{P}_a)$ according to Eq. F2 to obtain

$$\sum_a Q_a \bar{F}'(\tilde{P}_a) f_{r,a} > -c' d, \quad (\text{F7})$$

which is equivalent to the optimality condition at the boundary Eq. D10, provided that $\lambda^* = c' d$.

Appendix G: Cost function as a Lyapunov function of the dynamics

Here we show rigorously that, when the availability function is scale invariant, as in the case for the simple

cost function $F(m) = m^\alpha$, the dynamics must converge towards a fixed point. This fixed point is unique and corresponds to the optimal of the cost $\langle F \rangle$, as we have shown in the previous section.

$A(x)$ is scale invariant if there exists a function v such that $A(\gamma x) = v(\gamma)A(x)$. In this case we will see that the changes of relative frequencies P_r in the repertoire over time only depend on the total number of receptors through a prefactor. Below we derive the equations governing this dynamics and will then prove that this dynamics is assured to converge to a stable fixed point. We do so by showing that the dynamics admits the expected cost $\langle F \rangle$ as a Lyapunov function, *i.e.* a function that continually decreases under the dynamics.

For ease of notation we rewrite Eq. F1 as:

$$\frac{dN_r}{dt} = N_r[\pi_r(N) - d], \quad (\text{G1})$$

where N is a short-hand for $\{N_r\}$, and $\pi_r = \sum_a Q_a A(\sum_r N_r f_{r,a}) f_{r,a}$ is the growth rate of receptor type r . The relative frequencies $P_r = N_r/N_{\text{tot}}$ evolve according to:

$$\frac{dP_r}{dt} = \frac{1}{N_{\text{tot}}} \frac{dN_r}{dt} - \frac{N_r}{N_{\text{tot}}^2} \frac{dN_{\text{tot}}}{dt} \quad (\text{G2})$$

$$= P_r \left[\pi_r(N) - \sum_{r'} P_{r'} \pi_{r'}(N) \right]. \quad (\text{G3})$$

If A is scale invariant, so is π_r and $\pi_r(N) = \pi_r(N_{\text{tot}}P) = v(N_{\text{tot}})\pi_r(P)$. Then the equations further simplify to

$$\frac{dP_r}{dt} = v(N_{\text{tot}})P_r \left[\pi_r(P) - \sum_{r'} P_{r'} \pi_{r'}(P) \right], \quad (\text{G4})$$

$$= v(N_{\text{tot}})P_r (\pi_r - \bar{\pi}), \quad (\text{G5})$$

where $\bar{\pi} = \sum_r P_r \pi_r$.

We can now write how the expected cost $\langle F \rangle$ evolves in time:

$$\frac{d\langle F \rangle}{dt} = \sum_r \frac{\partial \langle F \rangle}{\partial P_r} \frac{dP_r}{dt} \quad (\text{G6})$$

$$= v(N_{\text{tot}}) \sum_r P_r \left[\sum_a Q_a \bar{F}'_a(\bar{P}_a) f_{r,a} \right] (\pi_r - \bar{\pi}) \quad (\text{G7})$$

$$= -\frac{v(N_{\text{tot}})}{c'} \sum_r P_r \left[\sum_a Q_a A(N_{\text{st}} \bar{P}_a) f_{r,a} \right] (\pi_r - \bar{\pi}) \quad (\text{G8})$$

$$= -\frac{v(N_{\text{tot}})v(N_{\text{st}})}{c'} \sum_r P_r \pi_r (\pi_r - \bar{\pi}) \quad (\text{G9})$$

$$= -\frac{v(N_{\text{tot}})v(N_{\text{st}})}{c'} \sum_r P_r (\pi_r - \bar{\pi})^2 \leq 0. \quad (\text{G10})$$

This proves that the cost always decreases with time, *i.e.* is a Lyapunov function of the dynamics. Therefore the dynamics will reach a stable fixed point at steady state, which is guaranteed to be the global minimum of the expected cost $\langle F \rangle$.

-
- [1] Burnet FM (1976) A modification of jerne's theory of antibody production using the concept of clonal selection. *The Australian Journal of Science* 20:67–69.
 - [2] Thomas N, et al. (2014) Tracking global changes induced in the cd4 t cell receptor repertoire by immunization with a complex antigen using short stretches of cdr3 protein sequence. *bioRxiv*.
 - [3] Weinstein JA, Jiang N, White RA, Fisher DS, Quake SR (2009) High-throughput sequencing of the zebrafish antibody repertoire. *Science* 324:807–810.
 - [4] Mora T, Walczak AM, Bialek W, Callan CG (2010) Maximum entropy models for antibody diversity. *Proc Natl Acad Sci USA* 107:5405–10.
 - [5] Ndifon W, et al. (2012) Chromatin conformation governs t-cell receptor j gene segment usage. *Proc Natl Acad Sci USA* 109:15865–70.
 - [6] Larimore K, McCormick MW, Robins HS, Greenberg PD (2012) Shaping of human germline igh repertoires revealed by deep sequencing. *J Immunol*.
 - [7] Sherwood AM, et al. (2011) Deep sequencing of the hu-

- man tcr and tcr repertoires suggests that tcr rearranges after and t cell commitment. *Sci Transl Med* 3:90ra61–90ra61.
- [8] Robins HS, et al. (2009) Comprehensive assessment of t-cell receptor beta-chain diversity in alphabeta t cells. *Blood* 114:4099–107.
- [9] Zvyagin IV, et al. (2014) Distinctive properties of identical twins' tcr repertoires revealed by high-throughput sequencing. *Proceedings of the National Academy of Sciences*.
- [10] Xia X (1998) How optimized is the translational machinery in escherichia coli, salmonella typhimurium and saccharomyces cerevisiae? *Genetics* 149:37–44.
- [11] Press WH (2009) Strong profiling is not mathematically optimal for discovering rare malfeasors. *Proceedings of the National Academy of Sciences* 106:1716–1719.
- [12] Perelson AS, Oster GF (1979) Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *Journal of theoretical biology* 81:645–670.

- [13] Chaikin PM, Lubensky TC (1995) *Principles of Condensed Matter Physics* (Cambridge University Press, Cambridge).
- [14] Torquato S, Stillinger FH (2003) Local density fluctuations, hyperuniformity, and order metrics. *Physical Review E* 68:041113.
- [15] Donev A, Stillinger FH, Torquato S (2005) Unexpected density fluctuations in jammed disordered sphere packings. *Phys Rev Lett* 95:090604.
- [16] Berthier L, Chaudhuri P, Coulais C, Dauchot O, Sollich P (2011) Suppressed compressibility at large scale in jammed packings of size-disperse spheres. *Phys Rev Lett* 106:120601.
- [17] Venturi V, Price DA, Douek DC, Davenport MP (2008) The molecular basis for public t-cell responses? *Nat Rev Immunol* 8:231–8.
- [18] De Boer RJ, Perelson AS (1994) T cell repertoires and competitive exclusion. *Journal of theoretical biology* 169:375–390.
- [19] De Boer RJ, Freitas AA, Perelson AS (2001) Resource competition determines selection of b cell repertoires. *Journal of theoretical biology* 212:333–343.
- [20] Vollmers C, Sit RV, Weinstein JA, Dekker CL, Quake SR (2013) Genetic measurement of memory b-cell recall using antibody repertoire sequencing. *Proc Natl Acad Sci USA* 110:13463–13468.
- [21] Moon JJ, et al. (2007) Naive cd4⁺ sup_h+i/sup_h t cell frequency varies for different epitopes and predicts repertoire diversity and response magnitude. *Immunity* 27:203–213.
- [22] Tkacik G, Walczak AM, Bialek W (2009) Optimizing information flow in small genetic networks. *Physical review E, Statistical, nonlinear, and soft matter physics* 80:031920.
- [23] Masland RH (1001) The fundamental plan of the retina. *Nature Neuroscience* 4:877–886.
- [24] Golisch T, Meister M (2010) Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron* 65:150–164.
- [25] Buck L, Axel R (1991) A novel multigene family may encode odorant receptors: A molecular basis for odor recognition. *Cell* 65:175–187.
- [26] Balasubramanian V, Sterling P (2009) Receptive fields and functional architecture in the retina. *The Journal of Physiology* 587:2753–2767.
- [27] MacArthur R, Levins R (1967) The limiting similarity, convergence, and divergence of coexisting species. *American naturalist* pp 377–385.
- [28] Szabó P, Meszéna G (2006) Limiting similarity revisited. *Oikos* 112:612–619.
- [29] Scheffer M, van Nes EH (2006) Self-organized similarity, the evolutionary emergence of groups of similar species. *Proceedings of the National Academy of Sciences* 103:6230–6235.
- [30] Pigolotti S, López C, Hernández-García E (2007) Species clustering in competitive lotka-volterra models. *Physical review letters* 98:258101.
- [31] Leimar O, Sasaki A, Doebeli M, Dieckmann U (2013) Limiting similarity, species packing, and the shape of competition kernels. *Journal of theoretical biology* 339:3–13.
- [32] Marraffini LA, Sontheimer EJ (2010) Crispr interference: Rna-directed adaptive immunity in bacteria and archaea. *Nat Rev Genet* 11:181–190.
- [33] He J, Deem MW (2010) Heterogeneous diversity of spacers within crispr (clustered regularly interspaced short palindromic repeats). *Phys Rev Lett* 105:128102.
- [34] Boyd SP, Vandenberghe L (2004) *Convex optimization* (Cambridge university press).
- [35] Bronshtein IN, Semendyayev KA, Musiol G, Muehlig H (2007) *Handbook of mathematics* (Springer) Vol. 3.
- [36] Parikh N, Boyd S (2013) Proximal algorithms. *Foundations and Trends in Optimization* 1:123–231.
- [37] Beck A, Teboulle M (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2:183–202.
- [38] Duchi J, Shalev-Shwartz S, Singer Y, Chandra T (2008) *Efficient projections onto the l 1-ball for learning in high dimensions* (ACM), pp 272–279.
- [39] Bertsekas DP (1999) *Nonlinear programming* (Athena Scientific).